

1 **Probabilistic precipitation type forecasting based on GEFS ensemble**
2 **forecasts of vertical temperature profiles**

3 Michael Scheuerer*, Scott Gregory

4 *University of Colorado, Cooperative Institute for Research in Environmental Sciences, and*
5 *NOAA/ESRL, Physical Sciences Division*

6 Thomas M. Hamill

7 *NOAA/ESRL, Physical Sciences Division*

8 Phillip E. Shafer

9 *NOAA/NWS/OST, Meteorological Development Laboratory*

10 **Corresponding author address:* Michael Scheuerer,

11 University of Colorado, Cooperative Institute for Research in Environmental Sciences, and

12 NOAA/ESRL, Physical Sciences Division 325 Broadway, R/PSD1, Boulder, CO 80305.

13 E-mail: michael.scheuerer@noaa.gov

ABSTRACT

14 A Bayesian classification method for probabilistic forecasts of precipitation
15 type is presented. The method considers the vertical wetbulb temperature pro-
16 files associated with each precipitation type, transforms them into their prin-
17 cipal components, and models each of these principal components by a skew
18 normal distribution. A variance inflation technique is used to de-emphasize
19 the impact of principal components corresponding to smaller eigenvalues,
20 and Bayes' theorem finally yields probability forecasts for each precipita-
21 tion type based on predicted wetbulb temperature profiles. Our approach is
22 demonstrated with reforecast data from the Global Ensemble Forecast Sys-
23 tem (GEFS) and observations at 551 METAR sites, using either the full en-
24 semble or the control run only. In both cases, reliable probability forecasts for
25 precipitation type being either rain, snow, ice pellets, freezing rain, or freez-
26 ing drizzle are obtained. Compared to the Model Output Statistics (MOS)
27 approach presently used by the National Weather Service, the skill of the pro-
28 posed method is comparable for rain and snow and significantly better for the
29 freezing precipitation types.

30 1. Introduction

31 Some forms of winter precipitation can have a substantial impact on air and ground transporta-
32 tion, and reliable predictions of them can help limit associated safety hazards and disruptions of
33 travel and commerce (Stewart et al. 2015, and references therein). Among several factors that
34 control the precipitation type at the surface, the vertical profile of wetbulb temperature T_w plays
35 a key role (e.g. Bourgozin 2000), and a number of algorithms have been devised which deter-
36 mine the precipitation type based on the T_w profile or quantities derived from it (e.g. Ramer 1993;
37 Baldwin et al. 1994; Bourgozin 2000; Schuur et al. 2012). A major challenge herein is the model
38 uncertainty about the T_w profile on the forecast day; while the above mentioned algorithms still
39 show good skill in detecting snow (SN) and rain (RA), reliable distinction between ice pellets
40 (IP) and freezing rain (FZRA) becomes increasingly difficult when this uncertainty is accounted
41 for (Reeves et al. 2014). A recently proposed algorithm, the spectral bin classifier (Reeves et al.
42 2016), pushes the limits of forecast accuracy for IP and FZRA by calculating the mass fraction of
43 liquid water for a spectrum of hydrometeors as they descend from the cloud top to the surface, thus
44 accounting for different rapidity of melting and refreezing of smaller hydrometeors compared to
45 larger ones. Their results still confirm the sensitivity of classification algorithms to perturbations
46 of the T_w profile. In a forecast setting where these profiles are derived from NWP model output
47 deviations of the true from the predicted (and interpolated) wetbulb temperature profile can be
48 substantial, especially for longer forecast lead times. In those situations with large uncertainty it
49 may be more useful to provide probability forecasts of each precipitation type, thus communicat-
50 ing the risk for precipitation to occur in the form of FZRA, say, instead of stating the most likely
51 outcome only. Operationally, such probabilistic guidance is currently provided for the contiguous
52 U.S. (CONUS) and Alaska by the Meteorological Development Laboratory (MDL) to support the

53 National Digital Guidance Database (NDGD). It is based on a model output statistics (MOS) ap-
54 proach that is described in Shafer (2010). This method links the probability of precipitation type
55 (PoPT) to NWP model output of variables such as 2-m temperature, 850 hPa temperature, 1000-
56 850 hPa thickness, 1000-500 hPa thickness, and freezing level. This approach yields conditional
57 probabilities of freezing (IP or FZRA), frozen (SN) and liquid (RA) precipitation for forecast lead
58 times up to 192 hours, but it does not attempt to distinguish the different freezing types. In this
59 paper we describe an alternative method which uses (discretized) vertical wetbulb temperature
60 profiles as a predictor, thus aiming to use more information from that profile as well as statisti-
61 cally modeling the forecast uncertainty. In Sec. 2 we describe the forecast and observation data
62 used in this study, which is identical to the data used by Shafer (2015), and which thus permits
63 a direct comparison between our approach and the operational method. Our statistical model and
64 the methods for fitting it to the training data are detailed in Sec. 3, while a detailed evaluation of
65 the precipitation type probabilities obtained with this model is the subject of Sec. 4. We finally
66 discuss the scope of our method and avenues for further improvement.

67 **2. Data used in this study**

68 *a. Observations*

69 Adopting the setup used by Shafer (2015), our method is calibrated with and validated against
70 weather observations at METAR (Meteorological Terminal Aviation Routine Weather Report) sites
71 (Allen and Erickson 2001a,b). Precipitation type observations were considered for the period
72 1996-2013 and all months between September and May (the period 09/1996 - 05/1997 will be
73 referred to as the '1996 cool season'), whenever precipitation was reported at the corresponding

74 site. Following Shafer (2015), we discarded sites where more than 50% of the precipitation type
75 reports were missing, leading to a set of 551 stations (506 CONUS, 26 Alaska, 19 Canada).

76 The original precipitation type reports, valid at 0000, 0600, 1200, and 1800 UTC, were classi-
77 fied into one of either three or five mutually exclusive categories. The first classification follows
78 the MOS precipitation type categories shown e.g. in Table 1 by Shafer (2015), and distinguishes
79 'freezing', 'frozen', and 'liquid' precipitation, classifying sleet as 'freezing' and any mixture of
80 liquid precipitation with snow as 'liquid' (Allen and Erickson 2001a,b; Shafer 2015). This three
81 category classification permits a direct comparison with the MOS technique used operationally by
82 the Meteorological Development Laboratory of the National Weather Service (NWS). In addition,
83 we consider a five-category classification which differs from the previous one in that it splits the
84 'freezing' category up into freezing rain (FZRA), freezing drizzle (FZDZ), and ice pellets (IP),
85 in order to study in how far our forecasts are able to provide probabilistic guidance that reflects
86 the new certification standards of the federal aviation administration (FAA) allowing some aircraft
87 to fly in FZDZ but not FZRA. The frozen and liquid types are relabeled as snow (SN) and rain
88 (RA) respectively. The observation data set used here clearly is not optimal for performing a five-
89 category classification. Only a subset of the 551 stations are augmented by human observers and
90 are able to report IP and FZDZ; all other stations may erroneously report a different type and thus
91 contaminate both training and verification samples with T_w profiles that should be associated with
92 IP or FZDZ but aren't. We accept the detrimental effect that this might have on the performance
93 of our method because our priority is a direct comparison with Shafer (2015), but we note that our
94 method might demonstrate somewhat better skill in distinguishing the different freezing precipi-
95 tation types if it were trained with an observation data set like the one from the mPING project
96 (Elmore et al. 2014) that is more consistent in how it reports IP, FZRA, and FZDZ.

97 *b. Forecasts*

98 Predictors used in this study were derived from the second-generation Global Ensemble Fore-
99 cast System (GEFS) reforecast data set (Hamill et al. 2013). GEFS data were extracted for 2-m
100 temperature and 2-m specific humidity on GEFS's native Gaussian grid at ~ 0.5 degree resolution
101 in an area surrounding the CONUS, Alaska, and southern Canada over the period 1996-2013. Sur-
102 face pressure and temperatures, specific humidities, and geopotential heights at the pressure levels
103 1000, 925, 850, 700, 500, and 300 hPa were obtained on a ~ 1 degree resolution grid covering the
104 same area. Horizontal grids were bilinearly interpolated to the station locations and were used to
105 convert the temperatures at the surface and at the pressure levels into wetbulb temperatures. Us-
106 ing the geopotential height fields, the wetbulb temperatures at each pressure level were associated
107 with a certain height above ground level (AGL) where ground level here refers to the GEFS model
108 grid. Vertical wetbulb temperature profiles were obtained by linear interpolation between those
109 pressure levels, and then discrete values were taken at fixed heights up to 3000 m AGL. As a result
110 of the rather coarse model grid resolution of ~ 0.5 degrees, the model grid elevation and the true
111 elevation at the station locations differ substantially in regions with complex terrain. In order to
112 adjust the resulting biases of the vertical wetbulb temperature profiles we:

- 113 • calculate the biases at the surface as the annual average difference between observed and
114 (horizontally interpolated) analyzed wetbulb temperatures at each location;
- 115 • assume that the bias linearly decreases to zero at the 500 hPa level;
- 116 • correct the entire vertical wetbulb temperature profile accordingly, i.e. apply the full (additive)
117 bias correction at the surface and gradually reduce the correction to zero with increasing
118 height above the ground.

119 This procedure is *not* meant to correct complex forecast biases which potentially have a seasonal
120 and diurnal cycle; these are somewhat implicitly addressed by the classification method described
121 in the subsequent section. The procedure described here only tries to remove biases resulting from
122 the mismatch between the terrain as represented by the forecast model and the true terrain.

123 For the remainder of this paper it is our working assumption that the observed precipitation type
124 only depends on the vertical wetbulb temperature profile above the ground, i.e. given two identical
125 profiles the outcome is independent of the location and time of the year at which those profiles were
126 observed. This is a simplification that does not account for the microphysical forcing (precipitation
127 rate, degree of riming, etc.) but is necessary as it allows us to pool data across all stations and
128 across all dates within the cool seasons considered here. Fig. 1 depicts examples of wetbulb
129 temperature profiles at initialization time (i.e. based on GFS analyses) obtained as described above.
130 There are typically only three or four pressure levels which data are used for the reconstruction
131 of the section of the profiles shown in this figure. It is clear that the resulting interpolation error
132 (compare e.g. with Fig. 2 in Reeves et al. 2014 who study wetbulb temperature profiles obtained
133 from radiosonde observations) can be substantial, and adds to the overall uncertainty about the
134 vertical profiles resulting from initial condition and forecast uncertainty.

135 **3. Regularized Bayesian Classification**

136 The method proposed in this paper is based on Bayes' Theorem, which has recently been em-
137 ployed by Hodyss et al. (2016) to derive optimal weights for different model forecasts and cli-
138 matology in a statistical postprocessing approach for continuous predictands. In our setting, the
139 predictand is categorical, but the same general principle can be used as a starting point. Assume
140 that we know, for each location s and each date t (day of the year and time of the day), the cli-
141 matological probability for each precipitation type $k \in \{1, \dots, K\}$ to occur. Denote this probability

142 by π_{kst} . Assume further that for each k we know the multivariate probability density function
 143 (PDF) ϕ_k that characterizes the distribution of the discretized, predicted vertical wetbulb temper-
 144 ature profiles that are compatible with the observed precipitation type k . For efficient statistical
 145 classification, we want these distributions to be as different as possible. Fig. 1 suggests that the
 146 differences between profiles corresponding to the different precipitation types are much more pro-
 147 nounced in the lower half of the sections depicted in the plots, and we therefore only consider
 148 wetbulb temperatures corresponding to heights above the surface up to 1500 m, even though the
 149 precipitation-generation layer is usually far outside this range. Sampling the profiles every 100 m
 150 then leaves us with a wetbulb temperature vector of dimension $d = 16$, and ϕ_k models the proba-
 151 bility distribution of this vector for each k . Due to our assumption that given two identical profiles
 152 the observed precipitation type should not depend on s and t , ϕ_k is assumed constant across the
 153 entire spatial domain and throughout the year (it may vary with lead time though since there is
 154 typically more dispersion around the mean profile for longer leads). According to Bayes' The-
 155 orem (Wilks 2006, Eq. (13.32)), given the climatological probabilities π_{kst} , the PDFs ϕ_k , and a
 156 new, predicted vector \mathbf{x} of vertical wetbulb temperature profile values, the conditional probability
 157 $P(k|\mathbf{x})$ of observing precipitation type k is

$$P(k|\mathbf{x}) = \frac{\pi_{kst} \phi_k(\mathbf{x})}{\sum_{i=1}^K \pi_{ist} \phi_i(\mathbf{x})}. \quad (1)$$

158 In this study we approximate the climatological probabilities π_{kst} by the relative frequencies of
 159 observed precipitation types, calculated separately for each location s , each month (but pooling
 160 all days within a month and all years for which data are available), and each time of the day.
 161 The following subsections discuss how an adequate model for ϕ_k can be defined and fitted. Note
 162 that Eq. (1) is also the starting point for (*quadratic*) *discriminant analysis*, where a deterministic
 163 classification rule is derived from this equation.

a. Basic model: multivariate normal distribution

A standard assumption with this approach to probabilistic classification is to let be ϕ_k a multivariate normal PDF (Wilks 2006, Sec. 13.3.3). This PDF is completely characterized by its mean vector μ_k and covariance matrix Σ_k . Given a set of training data we can estimate μ_k as the empirical mean and Σ_k as the empirical covariance matrix of the subset of the training profiles that correspond to an observed precipitation type k . In order to focus on situations where the outcome is truly uncertain, we only use locations/dates for the calculation of μ_k and Σ_k for which $\pi_{kst} < 0.99$ for every $k \in \{1, \dots, K\}$. This excludes, for example, precipitation events in January at high altitudes where precipitation most likely occurs in the form of snow, and precipitation events in May in Florida, where precipitation almost surely occurs in the form of rain. These events will still be used for validation, but excluding them for estimating the PDFs ϕ_k moves the mean vectors for rain and snow closer to the freezing point and improves the distribution fit in this temperature range where classification is most challenging. This results in a noticeable improvement in probabilistic classification skill, and one could even try to optimize the probability threshold for omitting cases from the training data set. However, we do not expect much further improvement from lowering the threshold and keep it fixed at 0.99.

Given μ_k, Σ_k , and a wetbulb temperature vector x , the likelihood $\phi_k(x)$ in Eq. (1) under the assumption of a multivariate normal distribution is given by

$$\phi_k(x) = (2\pi)^{-\frac{d}{2}} |\Sigma_k|^{-\frac{1}{2}} e^{-\frac{1}{2}(x-\mu_k)'\Sigma_k^{-1}(x-\mu_k)} \quad (2)$$

where x' denotes the transpose of x . Using the eigenvalue decomposition $\Sigma_k = E_k \Lambda_k E_k'$ to transform x into vectors $u_k = E_k'(x - \mu_k)$ of centered principal components (PCs), this likelihood

can be expressed as a product of univariate likelihoods

$$\varphi_k(\mathbf{x}) = \prod_{j=1}^d \phi_{0, \lambda_{k,j}}(u_{k,j}), \quad (3)$$

where $\phi_{(0, \lambda_{k,j})}$ is the PDF of a univariate normal distribution with mean 0 and variance equal to the j -th eigenvalue $\lambda_{k,j}$ of Σ_k . This re-interpretation of $\varphi_k(\mathbf{x})$ in terms of principal components will later be used to motivate an easily interpretable and computationally efficient generalization of the basic multivariate normal model discussed above. Fig. 2 shows the means of the $K = 5$ classes of interest and the variability around the respective mean in the direction of the first eigenvector of Σ_k . The different shapes of these eigenvectors suggest that there is structural information in the wetbulb temperature profiles beyond the mean that can be utilized for classification.

b. First extension: introducing skewness

Upon closer inspection, the assumption of a multivariate Gaussian distribution made above turns out to be a coarse approximation of the truth. For example, wetbulb temperature profiles much cooler than the mean profile are still compatible with observing snow, whereas the probability for observing snow but predicting a relatively warm profile decreases more rapidly (such profile would more be associated with observing rain). Applying a power transformation to each component of the wetbulb temperature vectors can make the distributions more symmetric (Wilks 2006, Sec. 3.4.1), but their direct physical interpretation is lost in that process. Alternatively, a more complex, multivariate skew normal distribution could be used to fit the untransformed data (Azzalini and Capitanio 1999). In our setting with dimension $d = 16$, this requires estimating a large number of model parameters which is computationally and numerically challenging. Here, we propose a similar approach that permits an intuitive interpretation and straightforward statistical inference. We first proceed as described above, estimate μ_k and Σ_k as the empirical means

and covariance matrices of the wetbulb temperature vectors associated with each precipitation type, and use them to calculate the centered PCs $u_{k,1}, \dots, u_{k,d}$ of each wetbulb temperature vector \mathbf{x} . Possible skewness can then be addressed for each PC separately by modeling them by univariate skew normal distributions $f_{\xi_{k,j}, \omega_{k,j}, \alpha_{k,j}}$ with location parameter $\xi_{k,j}$, scale parameter $\omega_{k,j}$, and shape parameter $\alpha_{k,j}$. By construction, the PCs are centered and have variances $\lambda_{k,j}$, so for given $\alpha_{k,j}$ the location and scale parameters are determined by

$$\omega_{k,j}^2 = \lambda_{k,j} \left(1 - \frac{2\alpha_{k,j}^2}{\pi(1 + \alpha_{k,j}^2)} \right)^{-1} \quad \text{and} \quad \xi_{k,j} = -\omega_{k,j} \sqrt{\frac{2\alpha_{k,j}^2}{\pi(1 + \alpha_{k,j}^2)}}. \quad (4)$$

Using these relations, $\alpha_{k,j}$ can be estimated via maximum likelihood. Since the impact of the PCs for smaller eigenvalues on classification will be de-emphasized as explained in the next subsection, we only bother to estimate $\alpha_{k,j}$ for $j \in \{1, 2\}$, and set $\alpha_{k,j} = 0$ (i.e. no skewness) for all other PCs. Fig. 3 shows histograms of the first three PCs and the fitted distributions. The variability in the direction of the first eigenvector corresponds to cooler/warmer than average wetbulb temperatures of the entire vertical profile (see Fig. 2), and the asymmetry of the associated PCs as described above for snow is clearly visible in those histograms. The fitted skew normal distributions are capable of modeling this asymmetry, and for calculating the likelihood $\phi_k(\mathbf{x})$ one only needs to replace Eq. (3) by

$$\phi_k(\mathbf{x}) = \prod_{j=1}^d f_{\xi_{k,j}, \omega_{k,j}, \alpha_{k,j}}(u_{k,j}). \quad (5)$$

c. Second extension: regularization

A further modification to the multivariate PDF ϕ_k is required to make this Bayesian classification method work efficiently. As pointed out in Sec. 2, the reconstruction of the vertical wetbulb temperature profile based on the GEFS model output at a few available pressure levels comes with substantial interpolation errors, and especially features at small vertical scales are not resolved.

225 On the other hand, even if we could reconstruct those profiles at high vertical resolution, it is
 226 unclear whether their fine-scale structure carries any useful information for the discrimination
 227 between different precipitation types. In the light of the principal component interpretation of the
 228 multivariate likelihood $\varphi_k(\mathbf{x})$ discussed above, it would seem natural to truncate after a few PCs
 229 and omit the last few terms in the product in Eq. (5), which typically correspond to eigenvectors
 230 representing the fine-scale structure of the vertical profiles. This is problematic, however, since the
 231 different precipitation types can have very different leading eigenvectors (see Fig. 2) and different
 232 spectra of eigenvalues, and so both the fractions of explained variances and the subspaces onto
 233 which the profiles are projected would be different. A more appropriate way to mute the effect of
 234 higher PCs on the likelihood $\varphi_k(\mathbf{x})$ is to *regularize* the covariance matrix Σ_k , i.e. to replace it by

$$\tilde{\Sigma}_k(a_k, b_k) := a_k \Sigma_k + b_k \mathbf{I}, \quad (6)$$

235 where \mathbf{I} is the identity matrix and a_k, b_k are positive coefficients for which selection will be dis-
 236 cussed later. This idea of regularization was introduced by Friedman (1989) in the context of a
 237 similar but deterministic classification technique referred to as *regularized discriminant analysis*.
 238 In our probabilistic setting we will refer to this idea as *regularized Bayesian classification (RBC)*.
 239 It can easily be combined with our assumption of skew normal distributions of the PCs by noting
 240 that the regularization in Eq. (6) leaves the eigenvectors unchanged but turns the eigenvalues $\lambda_{k,j}$
 241 into

$$\tilde{\lambda}_{k,j} = a_k \lambda_{k,j} + b_k, \quad j = 1, \dots, d. \quad (7)$$

242 Setting $a_k := 1 - b_k / \lambda_{k,1}$ leaves the first eigenvalue $\lambda_{k,1}$ unchanged but increases all other eigen-
 243 value with the relative increase being larger for smaller eigenvalues. This implies an artificial infla-
 244 tion of the variances of the univariate, skew normal PDFs in Eq. (5), and causes the corresponding
 245 likelihoods to be relatively less sensitive to the PCs $u_{k,j}$ corresponding to the smaller eigenvalues.

246 To find the optimal degree of inflation, i.e. optimal regularization parameters b_1, \dots, b_K , we use
 247 the training data set (forecasts and observations) that was used to estimate μ_k and Σ_k , and proceed
 248 as follows:

- 249 • for given parameters b_1, \dots, b_K , and for every wetbulb temperature profile \mathbf{x} in the training
 250 data set, use Eqs. (1), (5), (4), and (7) to calculate the likelihoods $\phi_k(\mathbf{x})$ and resulting forecast
 251 probabilities $P(k|\mathbf{x})$ for each k ; and
- 252 • use the corresponding training observations to calculate the resulting Brier skill scores BSS_k
 253 (see Sec. 4 for a definition) and choose b_1, \dots, b_K such that the sum $\sum_{k=1}^K \text{BSS}_k$ is maximized.

254 Note that the target function $\sum_{k=1}^K \text{BSS}_k$ that we seek to maximize gives the same weight to all
 255 precipitation type categories despite their very different frequencies of occurrence. This is done
 256 on purpose to foster good performance of our method with regard to the rare freezing precipitation
 257 types. Different priorities can be set, however, by introducing weights that increase or decrease
 258 the impact of the skill for certain precipitation types on the target function. In our example, the
 259 optimal values of b_k were between 5 and 10 for all k , which is smaller than the second, but about
 260 2-3 times larger than the third eigenvalues of Σ_k (see Fig. 3). This suggests that useful information
 261 about the vertical structure of the wetbulb temperature profiles is limited to the first two principal
 262 components.

263 *d. Third extension: applying RBC to ensemble forecasts*

264 The RBC approach presented above yields forecast probabilities for the occurrence of each
 265 precipitation type given a single (i.e. deterministic) forecast of a vertical wetbulb temperature
 266 profile. In our situation where we have an ensemble of forecasts, this ensemble represents some

267 of the uncertainty about the predicted vertical wetbulb temperature profiles, and its use can thus
 268 reduce the amount of variability that is modeled purely statistically.

269 We proceed as before regarding the estimation of μ_k , Σ_k , and the skewness parameters $\alpha_{k,j}$, con-
 270 sidering the forecast profiles x_1, \dots, x_M of the M ensemble members as separate cases. Centering
 271 and projecting those profiles onto the eigenvectors of Σ_k yields principal components $u_{k,j,m}$ and
 272 M different likelihoods $\varphi_k(x_m)$ for each class. The resulting probability forecasts $P(k|x_m)$ can be
 273 combined to a single probability forecast by simply taking the mean for each class

$$P(k|x_1, \dots, x_M) = \frac{1}{M} \sum_{m=1}^M P(k|x_m) \quad (8)$$

274 This way of linear pooling, however, does in general not yield reliable probability forecasts even
 275 if all of the individual member probability forecasts $P(k|x_m)$ are reliable (Ranjan and Gneiting
 276 2010). Indeed, as pointed out above, the simultaneous consideration of different ensemble member
 277 forecasts explains some of the variability of the forecast profiles, and so in return the statistically
 278 modeled variability needs to be reduced to avoid under-confident probability forecasts. We do this
 279 by replacing the variances $\lambda_{k,j}$ of the principal component PDFs that were originally obtained as
 280 the eigenvalues of Σ_k by the empirical variances $v_{k,j} = \text{var}(\bar{u}_{k,j})$ of the ensemble-mean PCs

$$\bar{u}_{k,j} = \frac{1}{M} \sum_{m=1}^M u_{k,j,m}.$$

281 While this is equivalent to just operating on the ensemble-mean profiles, evaluating Eq. (1) with
 282 an ensemble-mean profile does not yield the same probabilities as Eq. (8) due to the non-linearity
 283 of the likelihood function; Eq. (8) averages the probabilities corresponding to the different atmo-
 284 spheric situations represented by the ensemble, as opposed to averaging the vertical profiles and
 285 deriving a probability from the averaged state of the atmosphere. In contrast to $\lambda_{k,j}$, which de-
 286 scribes the variability of a certain PC over all profiles in the training data set corresponding to a
 287 certain precipitation type, $v_{k,j}$ elides the variability within the ensemble, and is therefore smaller

than $\lambda_{k,j}$. All subsequent steps, i.e. regularization according to Eq. (7), calculation of likelihoods according to Eqs. (5), (4), and calculation of probability forecasts according to Eqs. (1), (8), remain the same as before, but are carried out based on $v_{k,j}$ instead of $\lambda_{k,j}$. The results in the following section will show that this reduction of statistically modeled variability in favor of dynamically explained variability yields a noticeable improvement of predictive performance at longer forecast lead times.

4. Results

To test our RBC approach and compare it against the operational MOS technique, we adopt the verification setup used by Shafer (2015), but studying only the case of a training sample comprised of five cool seasons (see Sec. 2). Forecasts were produced and verified for the cool seasons 2001 through 2012. For each of these verification seasons, the methods were trained with data from the previous five cool seasons, i.e. the statistical model used for producing probability forecasts for the cool season 2001 was set up based on data (forecasts and observations) from the cool seasons 1996 through 2000. For estimating the climatological frequencies, which are used as a reference forecast on the one hand, and as a prior distribution for our RBC technique on the other hand, the entire observation record was used. This may be justified by noting that in practice long time series of observations are often available while forecast systems keep evolving and available forecast time series from a stable system are typically much shorter.

First, we assess the reliability of the RBC probability forecasts for different lead times separately for each precipitation type. Fig. 4 depicts reliability diagrams for probability forecasts generated by the ensemble-based version of the RBC method. For all forecast lead times considered in this study (including those not shown in this figure), the curves are close to the diagonal, which means

310 that the relative frequency of occurrence of each precipitation type matches the probability with
311 which it was predicted.

312 The RBC probability forecasts based on the GEFS control run only were equally reliable (not
313 shown here), so as a second validation tool we consider a quantitative performance measure, the
314 Brier skill score (Wilks 2006, Eqs. 7.34 and 7.35). In addition to reliability, the Brier score evalu-
315 ates the resolution of a forecast, i.e. its ability to distinguish situations with different frequencies
316 of occurrence. A skill score relates the score of the forecast method of interest to a reference score
317 (here: climatological frequency of occurrence, calculated separately for each location, each month,
318 and each time of the day) and thus facilitates its interpretation (Wilks 2006, Sec. 7.33). Here, we
319 compare the Brier skill scores (BSSs) of the generalized operator equation (GOE) implementation
320 of the operational MOS PoPT technique described in (Shafer 2010, 2015) and GEFS ensemble
321 mean forecasts, the RBC method based on the GEFS control run only, and the RBC method using
322 each of the individual ensemble member forecasts for forecast lead times up to 192h. The MOS
323 PoPT approach currently only distinguishes three classes: frozen (SN), liquid (RA), and freezing
324 (IP, FZRA or FZDZ). To allow a direct comparison, we aggregate the five class RBC probabili-
325 ties to three class probabilities, and compare the BSSs for the three class probabilities of all three
326 methods on the one hand, and the BSSs for the IP, FZRA and FZDZ probabilities by the two RBC
327 implementations on the other hand. The results depicted in Fig. 5 permit several conclusions:

- 328 • the use of ensemble forecasts as opposed to a single deterministic run clearly benefits forecast
329 performance, especially for longer forecast lead times;
- 330 • for the frozen and liquid class, the improvement of the RBC ensemble method over the MOS
331 PoPT approach is marginal; if the MOS PoPT approach were extended such as to use the

individual ensemble member forecasts rather than the ensemble mean, there may be no improvement at all;

- for the particularly challenging, freezing category, however, there is a noticeable benefit of using a statistical method (such as RBC) that can use the full vertical wetbulb temperature profile as a predictor; and
- the skill for the freezing categories (especially IP and FZDZ) is low compared to the skill for RA and SN; yet our RBC method can provide skillful probabilistic guidance on freezing precipitation several days ahead, and even has the potential to separate IP, FZRA and FZDZ.

The results discussed above show the effectiveness of our RBC method in general and the utility of ensemble forecasts in particular. How about the other two extensions (modeling skewness of the PCs, regularization), how much do they contribute to the skill of the RBC approach? How much skill is lost if the available training data for estimating μ_k , Σ_k , b_k , and $\alpha_{k,j}$ is composited of just one instead of five cool seasons? To answer these questions we use the control run based RBC method (fitted with five years of training data, as above) as a benchmark and compare it to a) the same model fitted with training data from a single cool season, b) a simplified model that regularizes Σ_k according to (6) but assumes normal instead of skew normal distributions of the PCs, c) a simplified model that uses skew normal distributions but does not regularize the empirical covariance matrices. The following conclusions can be drawn from the results shown in Fig. 6:

- a) Reducing the training sample size hardly affects the performance in predicting SN and RA probabilities, but has a rather strong, negative impact on the predictive performance for IP, FZRA, and FZDZ. For the two former, there are still enough cases within a single cool season to warrant a good estimation of model parameters. Estimating the parameters for the rare,

freezing precipitation types, however, requires either several years of training data or a much denser observation network. In addition to the issue of boundary discontinuity, this is also an argument in favor of a pooling data across all locations as opposed to partitioning the country into more homogeneous sub-domains. The latter might better account for different regional characteristics, but Fig. 6 suggests that these benefits could be nullified by the concomitant reduction of training sample size.

b) Simplifying the RBC approach by assuming a multivariate normal distribution for the wetbulb temperature profiles affects the predictive performance in the opposite way. While the more flexible distribution model does not seem to benefit the freezing precipitation types, the better approximation of the distributions of SN and RA profiles that results from modeling skewness in the PCs translates into improved skill of the resulting probability forecasts.

c) Finally, Fig. 6 highlights the necessity of regularizing the empirical covariance matrices. Without regularization, skill drops dramatically for SN and RA and becomes negative beyond a forecast lead time of three days. For the freezing types the impact is even stronger and lack of regularization results in Brier skill scores around -1.0 for all lead times. Unregularized classification gives as much emphasis to the noisy, unwarranted fine scale structure of the wetbulb temperature profiles as it gives to the first PCs that represent meaningful features of these profiles, and this results in probability forecasts that are entirely off the mark.

To illustrate the capabilities and limits of probabilistic guidance obtained with the RBC method applied to GEFS ensemble forecasts, two particular cases studies presented. Fig. 7 shows spatial maps of FZRA probabilities for 27 January 2009, 0000 UTC, with a forecast lead time of 2, 4, and 6 days ahead. This date is in the middle of a major ice storm that impacted parts of Oklahoma, Arkansas, Missouri, Illinois, Indiana, West Virginia, and Kentucky. The plots suggest that

the GEFS captured the atmospheric situation well, and the RBC methods provides a strong probabilistic signal for freezing rain even at 6 days of lead time. For the event shown in Fig. 8 (also studied by Reeves et al. 2016) the situation is more complex. The plots show observed precipitation types and 2 day ahead RBC forecast probabilities for 22 February 2013, 0000 UTC. Even at this short lead time, the probabilistic signal for the freezing precipitation types is rather weak (note the different color scales) and no clear guidance is provided as to which particular freezing precipitation type will dominate in each geographical area. This underscores the inherent uncertainty in precipitation type forecasts based on a global ensemble prediction system, and illustrates the limits of such forecasts. Notwithstanding, the RBC probability forecasts indicate an increased risk of freezing precipitation, and we believe that there is substantial value in communicating that risk to decision makers.

5. Discussion

In this paper we have proposed a method for conditional probabilistic precipitation type forecasting which is based on a statistical model for the predicted vertical wetbulb temperature profiles that are compatible with each precipitation type. Using Bayes' theorem this model can be inverted such that it yields probability forecasts for each precipitation type given a new predicted profile.

There were many sources of forecast and data uncertainty that needed to be accounted for in a precipitation typing methodology. These include forecast errors stemming from initial condition uncertainty, from model error, and in this case from the need to interpolate NWP model output from a relatively coarse horizontal grid and a few pressure levels to a much finer horizontal and vertical resolution and more complex orography at the surface level. Availability of sigma-level forecast data at a finer vertical resolution could reduce this last component of uncertainty, which contributes noticeably to the overall uncertainty about the wetbulb temperature profiles at short

401 lead times. It is suggested that thermodynamic variables be archived at many vertical levels above
402 the surface when generating future reforecasts. At longer lead times, forecast errors become the
403 dominant source of uncertainty, and the interpolation error might be negligible. At short lead
404 times, forecasts from a high resolution, limited-area NWP model might be available, which might
405 be accurate enough to yield superior classification results using an explicit precipitation type di-
406 agnosis scheme (e.g. Benjamin et al. 2016) or the spectral bin classifier proposed by Reeves et al.
407 (2016), but such guidance could be leveraged in a probabilistic framework, too.

408 The strength of the method proposed here is that it can handle the large uncertainty that in-
409 evitably comes with predictions from a global forecast system, and that it can still provide reli-
410 able, probabilistic precipitation type forecasts at forecast lead times up to seven days ahead. It has
411 sufficient skill to give decision makers at least a heads up about precipitation type related weather
412 risks, and it can easily be extended to distinguish further precipitation type classes like mixtures
413 of snow and rain, mixtures of freezing precipitation types, and so forth, if they are reported accu-
414 rately in the observations. The observation data set used here is not optimal in that regard as it is
415 inconsistent in how it reports IP, FZRA, and FZDZ, and the skill of our method in distinguishing
416 these types might actually be better than reported here if it were trained with an observation data
417 set like the one from the mPING project (Elmore et al. 2014, 2015) in which IP, FZRA, and FZDZ
418 are distinguished more systematically.

419 We have focused on vertical profiles of wetbulb temperature as a predictor variable. However,
420 by combining the statistical dimension reduction / regularization techniques used here with more
421 physically motivated aggregation methods one might be able to further improve skill by using ad-
422 ditional predictors such as relative humidity profiles. Alternatively, one could use modern machine
423 learning techniques like neural networks to identify features of vertical wetbulb temperature and
424 humidity profiles that determine the observed precipitation type. While extremely powerful, these

425 techniques typically require large data sets for training, but these may become available once sev-
426 eral years of mPING data have been collected, and allow one to explore the more data-intensive
427 machine learning techniques.

428 *Acknowledgments.* The authors thank Amanda Hering for useful discussions which inspired
429 some of the refinements to our basic distribution model for the wetbulb temperature profiles. Our
430 research was supported by grants from the NOAA/NWS Sandy Supplemental (Disaster Relief
431 Appropriations Act of 2013) and the NOAA/NWS Research to Operations (R2O) initiative for the
432 Next-Generation Global Prediction System (NGGPS), award # NA15OAR4320137.

433 **References**

434 Allen, R. L., and M. C. Erickson, 2001a: AVN-based MOS precipitation type guidance for the
435 United States. NWS Technical Procedures Bulletin No. 476, NOAA, U.S. Dept. of Commerce,
436 8 pp.

437 Allen, R. L., and M. C. Erickson, 2001b: MRF-based MOS precipitation type guidance for the
438 United States. NWS Technical Procedures Bulletin No. 485, NOAA, U.S. Dept. of Commerce,
439 8 pp.

440 Azzalini, A., and A. Capitanio, 1999: Statistical applications of the multivariate skew normal
441 distribution. *J. Roy. Stat. Soc. B*, **61**, 579–602.

442 Baldwin, M., R. Treadon, and S. Contorno, 1994: Precipitation type prediction using a decision
443 tree approach with NMCs mesoscale eta model. *10th Conf. on Numerical Weather Prediction*,
444 Portland, OR, Amer. Meteor. Soc., 30–31.

Benjamin, S. G., J. M. Brown, and T. G. Smirnova, 2016: Explicit precipitation-type diagnosis from a model using a mixed-phase bulk cloud-precipitation microphysics parameterization. *Wea. Forecasting*, **31**, 609–619.

Bourgouin, P., 2000: A method to determine precipitation type. *Wea. Forecasting*, **15**, 583–592.

Elmore, K. L., Z. L. Flamig, V. Lakshmanan, B. T. Kaney, V. Farmer, H. D. Reeves, and L. S. Rothfusz, 2014: mPING: Crowd-sourcing weather reports for research. *Bull. Amer. Meteor. Soc.*, **95**, 1335–1342.

Elmore, K. L., H. M. Grams, D. Apps, and H. D. Reeves, 2015: Verifying forecast precipitation type with mPING. *Wea. Forecasting*, **30**, 656–667.

Friedman, J. H., 1989: Regularized discriminant analysis. *J. Amer. Stat. Assoc.*, **84**, 165–175.

Hamill, T. M., G. T. Bates, J. S. Whitaker, D. R. Murray, M. Fiorino, T. J. G. Jr., Y. Zhu, and W. Lapenta, 2013: NOAA’s second-generation global medium-range ensemble reforecast data set. *Bull. Amer. Meteor. Soc.*, **94**, 1553–1565.

Hodyss, D., E. Satterfield, J. McLay, T. M. Hamill, and M. Scheuerer, 2016: Inaccuracies with multimodel postprocessing methods involving weighted, regression-corrected forecasts. *Mon. Wea. Rev.*, **144**, 1649–1668.

Ramer, J., 1993: An empirical technique for diagnosing precipitation type from model output. *Fifth Int. Conf. on Aviation Weather Systems*, Vienna, VA, Amer. Meteor. Soc., 227–230.

Ranjan, R., and T. Gneiting, 2010: Combining probability forecasts. *J. Roy. Stat. Soc. B*, **32**, 71–91.

Reeves, H. D., K. L. Elmore, A. V. Ryzhkov, T. J. Schuur, and J. Krause, 2014: Sources of uncertainty in precipitation-type forecasting. *Wea. Forecasting*, **29**, 936–953.

- 467 Reeves, H. D., A. V. Ryzhkof, and J. Krause, 2016: Discrimination between winter precipitation
468 types based on spectral-bin microphysical modeling. *J. Appl. Meteor. Climatol.*, **55**, 1747–1761.
- 469 Schuur, T. J., H.-S. Park, A. V. Ryzhkof, and H. D. Reeves, 2012: Classification of precipitation
470 types during transitional winter weather using the ruc model and polarimetric radar retrievals.
471 *J. Appl. Meteor. Climatol.*, **51**, 763–779.
- 472 Shafer, P., 2010: Logit transforms in forecasting precipitation type. *20th Conf. on Probability and*
473 *Statistics in the Atmospheric Sciences*, Atlanta, GA, Amer. Meteor. Soc., P222.
- 474 Shafer, P. E., 2015: A sample size sensitivity test for MOS precipitation type. *Special Symposium*
475 *on Model Postprocessing and Downscaling*, Phoenix, AZ, Amer. Meteor. Soc., 4.1.
- 476 Stewart, R. E., J. M. Thériault, and W. Henson, 2015: On the characteristics of and processes
477 producing winter precipitation types near 0°C. *Bull. Amer. Meteor. Soc.*, **96**, 623–639.
- 478 Wilks, D. S., 2006: *Statistical Methods in the Atmospheric Sciences*, International Geophysics
479 Series, Vol. 91. 2nd ed., Elsevier Academic Press.

LIST OF FIGURES

Fig. 1.	Approximate vertical wetbulb temperature profiles reconstructed from GEFS forecasts fields at initialization time. For each of the five precipitation types of interest, 30 profiles are depicted that were randomly sampled from locations/dates where that precipitation type was reported.	25
Fig. 2.	Empirical means (solid lines) and variability (two standard deviations, dashed lines) in the direction of the first eigenvector of Σ_k for each of the five precipitation types distinguished by our algorithm.	26
Fig. 3.	Histograms and fitted skew normal distributions for the first three principal components of the wetbulb temperature profile vectors of each class.	27
Fig. 4.	Reliability diagrams for probability forecasts at various forecast lead times generated by the ensemble-based version of the RBC method with probabilities rounded to a precision of 0.05. The inset histograms depict the frequency (on a logarithmic scale) with which the respective probabilities were forecast (x-axes are the same as in the reliability diagrams). Points of the reliability curve associated with very infrequent forecast probabilities (< 25 cases) are subject to substantial sampling variability and have therefore been omitted.	28
Fig. 5.	Brier skill scores of the probability forecasts obtained with the operational MOS approach and the RBC methods based on the control run and based on the full ensemble.	29
Fig. 6.	Brier skill scores of the probability forecasts obtained with variants of the RBC method (applied to the GEFS control forecast) which use only one year of training data, normal instead of skew normal distributions, or unregularized covariance matrices Σ_k	30
Fig. 7.	Observed precipitation types (shaded circles with color scheme as in previous figures) on 27 January 2009, 0000 UTC, and predicted freezing rain probabilities by the ensemble-based RBC method for different forecast lead times.	31
Fig. 8.	Observed precipitation types on 22 February 2013, 0000 UTC, and predicted precipitation type probabilities by the ensemble-based RBC method with a forecast lead time of 48h.	32




FIG. 1. Approximate vertical wetbulb temperature profiles reconstructed from GEFS forecasts fields at ini-
tialization time. For each of the five precipitation types of interest, 30 profiles are depicted that were randomly
sampled from locations/dates where that precipitation type was reported.

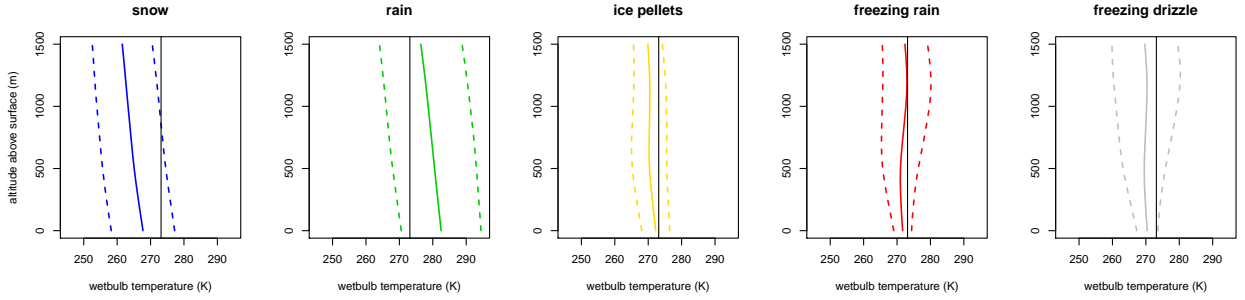


FIG. 2. Empirical means (solid lines) and variability (two standard deviations, dashed lines) in the direction
of the first eigenvector of Σ_k for each of the five precipitation types distinguished by our algorithm.

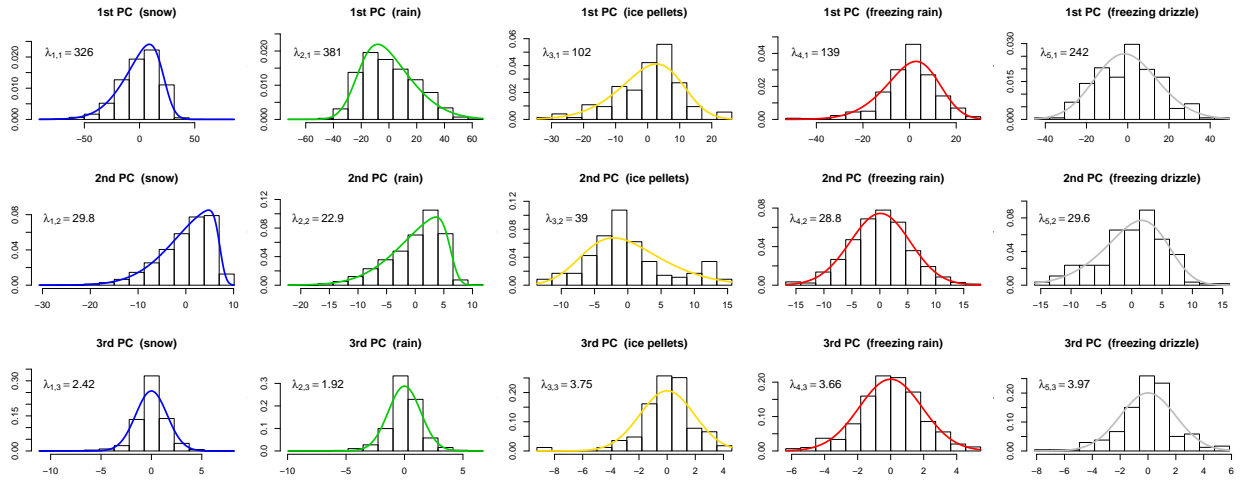


FIG. 3. Histograms and fitted skew normal distributions for the first three principal components of the wetbulb temperature profile vectors of each class.

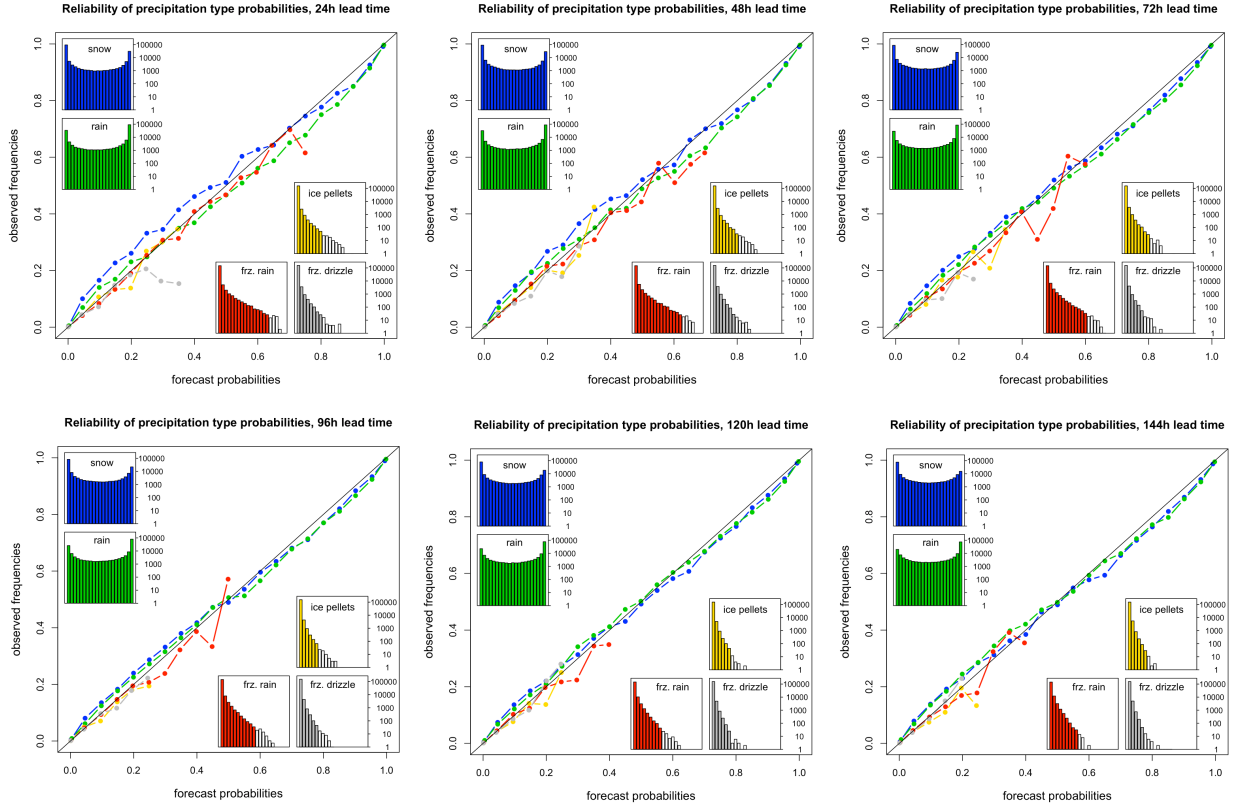


FIG. 4. Reliability diagrams for probability forecasts at various forecast lead times generated by the ensemble-based version of the RBC method with probabilities rounded to a precision of 0.05. The inset histograms depict the frequency (on a logarithmic scale) with which the respective probabilities were forecast (x-axes are the same as in the reliability diagrams). Points of the reliability curve associated with very infrequent forecast probabilities (< 25 cases) are subject to substantial sampling variability and have therefore been omitted.

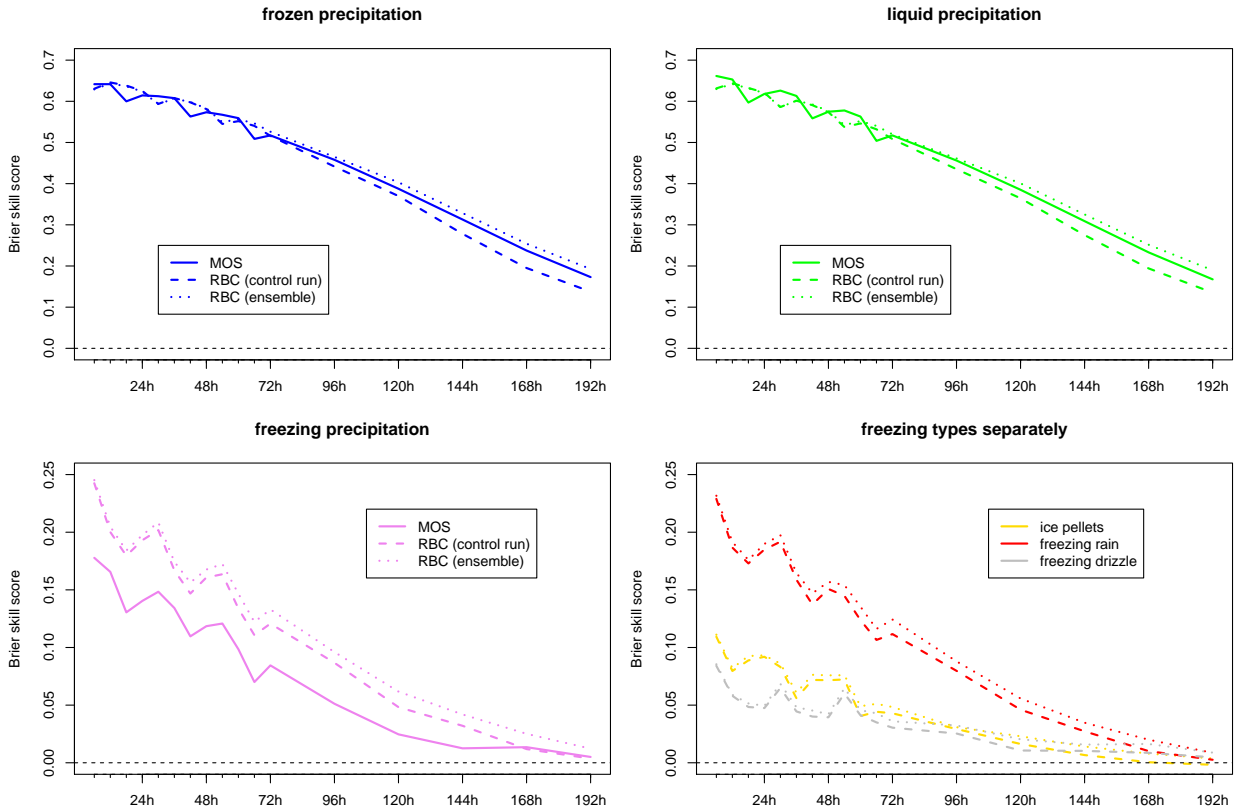


FIG. 5. Brier skill scores of the probability forecasts obtained with the operational MOS approach and the RBC methods based on the control run and based on the full ensemble.

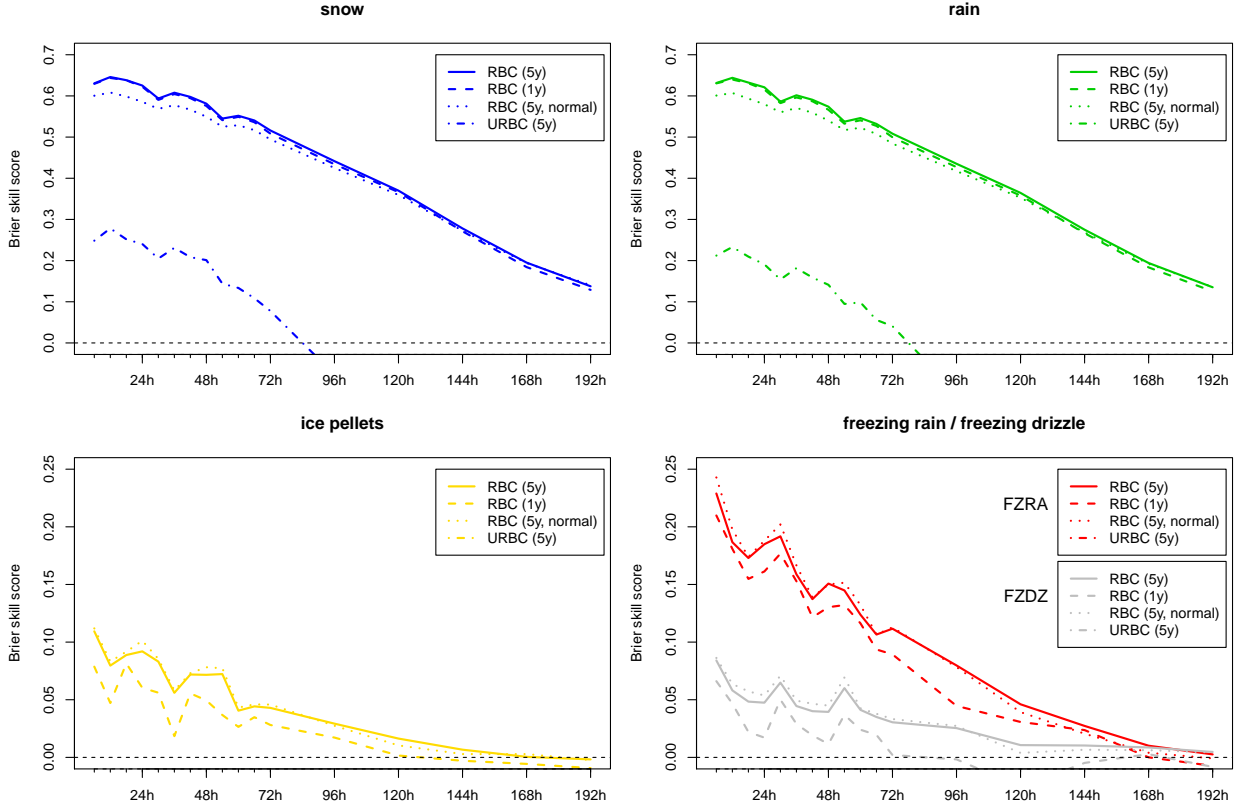


FIG. 6. Brier skill scores of the probability forecasts obtained with variants of the RBC method (applied to the GEFS control forecast) which use only one year of training data, normal instead of skew normal distributions, or unregularized covariance matrices Σ_k .

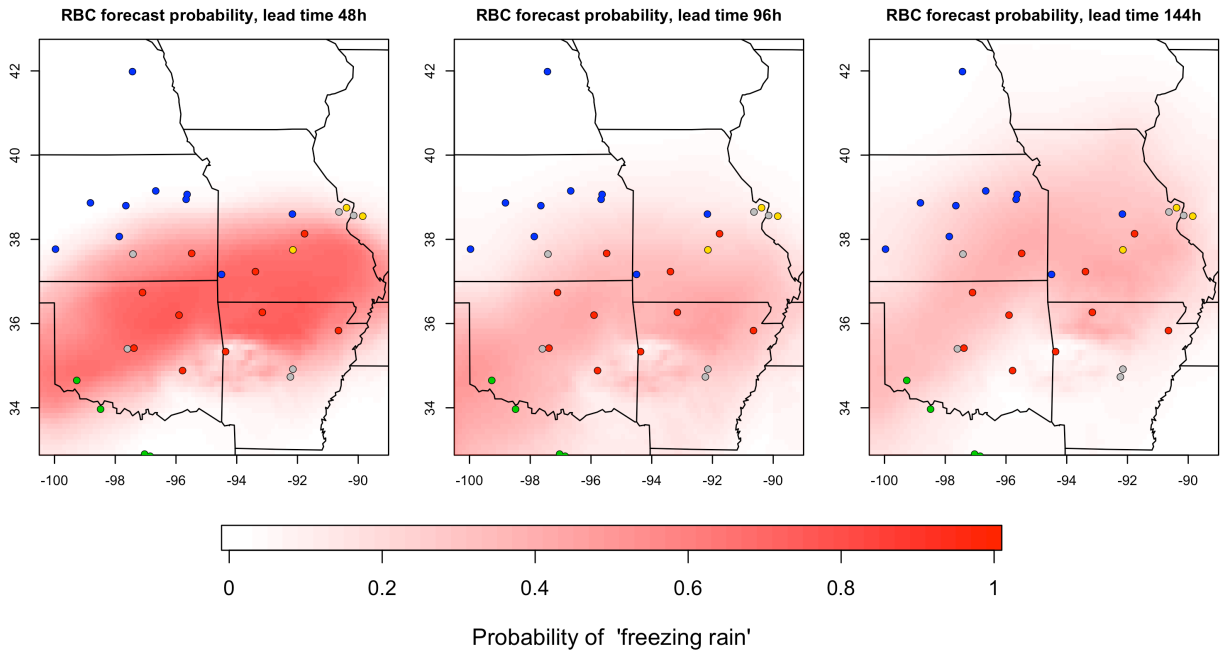


FIG. 7. Observed precipitation types (shaded circles with color scheme as in previous figures) on 27 January 2009, 0000 UTC, and predicted freezing rain probabilities by the ensemble-based RBC method for different forecast lead times.

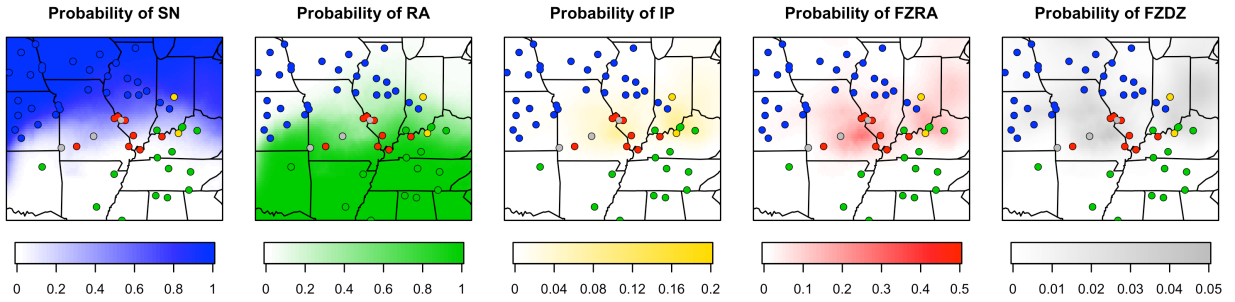


FIG. 8. Observed precipitation types on 22 February 2013, 0000 UTC, and predicted precipitation type probabilities by the ensemble-based RBC method with a forecast lead time of 48h.